



[> home](#) [> about](#) [> feedback](#) [> login](#)

USPTO



Try the *new* Portal design

Give us your opinion after using it.

## Search Results

Search Results for: **[topology <and> cache <and> query <and> plan<AND> ((inverted <and> indice <and> query) )]**  
Found **21** of **154,226** searched.

## Search within Results



[> Advanced Search](#)

[> Search Help/Tips](#)

Sort by: **Title** **Publication** **Publication Date** **Score**  **Binder**


Results 1 - 20 of 21 short listing


  
Prev  
Page

1 2

  
Next  
Page

- 1  Distributed indexing: a scalable mechanism for distributed information retrieval 82%  
Peter B. Danzig , Jongsuk Ahn , John Noll , Katia Obraczka  
**Proceedings of the 14th annual international ACM SIGIR conference on Research and development in information retrieval** September 1991

- 2  ScentTrails: Integrating browsing and searching on the Web 80%  
Christopher Olston , Ed H. Chi  
**ACM Transactions on Computer-Human Interaction (TOCHI)** September 2003  
Volume 10 Issue 3  
The two predominant paradigms for finding information on the Web are browsing and keyword searching. While they exhibit complementary advantages, neither paradigm alone is adequate for complex information goals that lend themselves partially to browsing and partially to searching. To integrate browsing and searching smoothly into a single interface, we introduce a novel approach called ScentTrails. Based on the concept of information scent developed in the context of information foraging theory,  
...

- 3  External memory algorithms and data structures: dealing with **massive data** 80%  
Jeffrey Scott Vitter  
**ACM Computing Surveys (CSUR)** June 2001  
Volume 33 Issue 2  
Data sets in large applications are often too massive to fit completely inside the computers internal memory. The resulting input/output communication (or I/O) between fast internal memory and slower external memory (such as disks) can be a major performance bottleneck. In this article we survey the state of the art in the

design and analysis of external memory (or EM) algorithms and data structures, where the goal is to exploit locality in order to reduce the I/O costs. We consider a varie ...

#### 4 Version models for software configuration management 80%



Reidar Conradi , Bernhard Westfechtel

**ACM Computing Surveys (CSUR)** June 1998

Volume 30 Issue 2

After more than 20 years of research and practice in software configuration management (SCM), constructing consistent configurations of versioned software products still remains a challenge. This article focuses on the version models underlying both commercial systems and research prototypes. It provides an overview and classification of different versioning paradigms and defines and relates fundamental concepts such as revisions, variants, configurations, and changes. In particular, we foc ...

#### 5 An Experimental Study of Polylogarithmic, Fully Dynamic, Connectivity Algorithms 80%



Raj Iyer , David Karger , Hariharan Rahul , Mikkel Thorup

**Journal of Experimental Algorithmics (JEA)** January 2001

Volume 6

We present an experimental study of different variants of the amortized  $O(\log^2 n)$ -time fully-dynamic connectivity algorithm of Holm, de Lichtenberg, and Thorup (STOC'98). The experiments build upon experiments provided by Alberts, Cattaneo, and Italiano (SODA'96) on the randomized amortized  $O(\log^3 n)$  fully-dynamic connectivity algorithm of Henzinger and King (STOC'95). Our experiments shed light upon similarities and differences between ...

#### 6 Testbed directions and experience: PlanetLab: an overlay testbed for broad-coverage services 77%



Brent Chun , David Culler , Timothy Roscoe , Andy Bavier , Larry Peterson , Mike Wawrzoniak , Mic Bowman

**ACM SIGCOMM Computer Communication Review** July 2003

Volume 33 Issue 3

PlanetLab is a global overlay network for developing and accessing broad-coverage network services. Our goal is to grow to 1000 geographically distributed nodes, connected by a diverse collection of links. PlanetLab allows multiple service to run concurrently and continuously, each in its own slice of PlanetLab. This paper describes our initial implementation of PlanetLab, including the mechanisms used to implement virtualization, and the collection of core services used to manage PlanetLab.

#### 7 Special topic section on peer to peer data management: Toward network data independence 77%



Joseph M. Hellerstein

**ACM SIGMOD Record** September 2003

Volume 32 Issue 3


A number of researchers have become interested in the design of global-scale networked systems and applications. Our thesis here is that the database community's principles and technologies have an important role to play in the design of these systems. The point of departure is at the roots of database research: we generalize Codd's notion of data independence to physical environments beyond storage systems. We note analogies between the development of database indexes and the new generation of ...

## 8 Peer-to-peer: Making gnutella-like P2P systems scalable 77%

 Yatin Chawathe , Sylvia Ratnasamy , Lee Breslau , Nick Lanham , Scott Shenker  
**Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications** August 2003

Napster pioneered the idea of peer-to-peer file sharing, and supported it with a centralized file search facility. Subsequent P2P systems like Gnutella adopted decentralized search algorithms. However, Gnutella's notoriously poor scaling led some to propose distributed hash table solutions to the wide-area file search problem. Contrary to that trend, we advocate retaining Gnutella's simplicity while proposing new mechanisms that greatly improve its scalability. Building upon prior research [1, 1 ...

## 9 Decentralized storage systems: Farsite: federated, available, and 77%

 reliable storage for an incompletely trusted environment  
 Atul Adya , William J. Bolosky , Miguel Castro , Gerald Cermak , Ronnie Chaiken , John R. Douceur , Jon Howell , Jacob R. Lorch , Marvin Theimer , Roger P. Wattenhofer  
**ACM SIGOPS Operating Systems Review** December 2002  
 Volume 36 Issue SI


Farsite is a secure, scalable file system that logically functions as a centralized file server but is physically distributed among a set of untrusted computers. Farsite provides file availability and reliability through randomized replicated storage; it ensures the secrecy of file contents with cryptographic techniques; it maintains the integrity of file and directory data with a Byzantine-fault-tolerant protocol; it is designed to be scalable by using a distributed hint mechanism and delegatio ...

## 10 Dynamic services and analysis: Make it fresh, make it quick: searching 77%

 a network of personal webserver  
 Mayank Bawa , Roberto J. Bayardo , Sridhar Rajagopalan , Eugene J. Shekita  
**Proceedings of the twelfth international conference on World Wide Web** May 2003

Personal webserver have proven to be a popular means of sharing files and peer collaboration. Unfortunately, the transient availability and rapidly evolving content on such hosts render centralized, crawl-based search indices stale and incomplete. To address this problem, we propose YouSearch, a distributed search application for personal webserver operating within a shared context (e.g., a corporate intranet). With YouSearch, search results are always fast, fresh and complete -- properties we ...

## 11 Database indexing for large DNA and protein sequence collections 77%

 Ela Hunt , Malcolm P. Atkinson , Robert W. Irving  
**The VLDB Journal – The International Journal on Very Large Data Bases**  
 November 2002  
 Volume 11 Issue 3

Our aim is to develop new database technologies for the approximate matching of unstructured string data using indexes. We explore the potential of the suffix tree data structure in this context. We present a new method of building suffix trees, allowing us to build trees in excess of RAM size, which has hitherto not been possible. We show that this method performs in practice as well as the  $O(n)$  method of Ukkonen [70]. Using this method we build indexes for 200 Mb of protein and 3 ...

## 12 XML schemas: integration and translation: A local search mechanism for 77%

 peer-to-peer networks  
 Vana Kalogeraki , Dimitrios Gunopulos , D. Zeinalipour-Yazti  
**Proceedings of the eleventh international conference on Information and**

**knowledge management** November 2002

One important problem in peer-to-peer (P2P) networks is searching and retrieving the correct information. However, existing searching mechanisms in pure peer-to-peer networks are inefficient due to the decentralized nature of such networks. We propose two mechanisms for information retrieval in pure peer-to-peer networks. The first, the modified Breadth-First Search (BFS) mechanism, is an extension of the current Gnutella protocol, allows searching with keywords, and is designed to minimize the ...

**13 A case for dynamic view management**

77%



Yannis Kotidis , Nick Roussopoulos

**ACM Transactions on Database Systems (TODS)** December 2001

Volume 26 Issue 4

Materialized aggregate views represent a set of redundant entities in a data warehouse that are frequently used to accelerate On-Line Analytical Processing (OLAP). Due to the complex structure of the data warehouse and the different profiles of the users who submit queries, there is need for tools that will automate and ease the view selection and management processes. In this article we present DynaMat, a system that manages dynamic collections of materialized aggregate views in a data warehouse ...

**14 A tool for Internet-oriented knowledge based systems**

77%



Robert Inder

**Proceedings of the 2000 ACM symposium on Applied computing** March 2000**15 Piranha: a scalable architecture based on single-chip multiprocessing**

77%



Luiz André Barroso , Kourosh Gharachorloo , Robert McNamara , Andreas Nowatzky , Shaz Qadeer , Barton Sano , Scott Smith , Robert Stets , Ben Verghese

**ACM SIGARCH Computer Architecture News , Proceedings of the 27th annual international symposium on Computer architecture** May 2000

Volume 28 Issue 2

The microprocessor industry is currently struggling with higher development costs and longer design times that arise from exceedingly complex processors that are pushing the limits of instruction-level parallelism. Meanwhile, such designs are especially ill suited for important commercial applications, such as on-line transaction processing (OLTP), which suffer from large memory stall times and exhibit little instruction-level parallelism. Given that commercial applications constitute by fa ...

**16 On network-aware clustering of Web clients**

77%



Balachander Krishnamurthy , Jia Wang

**ACM SIGCOMM Computer Communication Review , Proceedings of the conference on Applications, Technologies, Architectures, and Protocols for Computer Communication** August 2000

Volume 30 Issue 4

Being able to identify the groups of clients that are responsible for a significant portion of a Web site's requests can be helpful to both the Web site and the clients. In a Web application, it is beneficial to move content closer to groups of clients that are responsible for large subsets of requests to an origin server. We introduce clusters---a grouping of clients that are close together topologically and likely to be under common administrative control. We identify clu ...

**17 A case for intelligent disks (IDISks)**

77%



Kimberly Keeton , David A. Patterson , Joseph M. Hellerstein

**ACM SIGMOD Record** September 1998

## Volume 27 Issue 3

Decision support systems (DSS) and data warehousing workloads comprise an increasing fraction of the database market today. I/O capacity and associated processing requirements for DSS workloads are increasing at a rapid rate, doubling roughly every nine to twelve months [38]. In response to this increasing storage and computational demand, we present a computer architecture for decision support database servers that utilizes "intelligent" disks (IDISks). IDISks utilize low-cost ...

**18** External memory algorithms

77%



Jeffrey Scott Vitter

**Proceedings of the seventeenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems** May 1998

**19** Technique for automatically correcting words in text

77%



Karen Kukich

**ACM Computing Surveys (CSUR)** December 1992

Volume 24 Issue 4

Research aimed at correcting words in text has focused on three progressively more difficult problems: (1) nonword error detection; (2) isolated-word error correction; and (3) context-dependent word correction. In response to the first problem, efficient pattern-matching and n-gram analysis techniques have been developed for detecting strings that do not appear in a given word list. In response to the second problem, a variety of general and application-specific spelling cor ...

**20** Query evaluation techniques for large databases

77%



Goetz Graefe



**ACM Computing Surveys (CSUR)** June 1993

Volume 25 Issue 2

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

---

**Results 1 - 20 of 21      short listing**


  
 Prev Page   1   2   Next Page

---

 The ACM Portal is published by the Association for Computing Machinery. Copyright ?2005 ACM, Inc.



[> home](#) [> about](#) [> feedback](#) [> login](#)

USPTO



Try the *new* Portal design

Give us your opinion after using it.

## Search Results

Search Results for: **[topology <and> cache <and> query <and> plan <and> evalution ]**

Found **1** of **154,226** searched.

## Search within Results



[> Advanced Search](#)

[> Search Help/Tips](#)

---

**Sort by:** [Title](#) [Publication](#) [Publication Date](#) [Score](#)  [Binder](#)

---

**Results 1 - 1 of 1**    [short listing](#)

---

**1** Special issue on knowledge representation 77%



Ronald J. Brachman , Brian C. Smith

**ACM SIGART Bulletin** February 1980

Issue 70

In the fall of 1978 we decided to produce a special issue of the SIGART Newsletter devoted to a survey of current knowledge representation research. We felt that there were two useful functions such an issue could serve. First, we hoped to elicit a clear picture of how people working in this subdiscipline understand knowledge representation research, to illuminate the issues on which current research is focused, and to catalogue what approaches and techniques are currently being developed.  
Secon ...

---

**Results 1 - 1 of 1**    [short listing](#)

---

The ACM Portal is published by the Association for Computing Machinery. Copyright ?2005 ACM, Inc.



> [home](#) > [about](#) > [feedback](#) > [login](#)

USPTO



Try the *new* Portal design

Give us your opinion after using it.

## Search Results

Search Results for: **[topology <and> cache <and> query <and> plan <and> directory <and> server <and> distribute <and> forest]**

Found **8** of **154,226** searched.

## Search within Results



> [Advanced Search](#)

> [Search Help/Tips](#)

**Sort by:** [Title](#) [Publication](#) [Publication Date](#) [Score](#)  [Binder](#)

**Results 1 - 8 of 8**    [short listing](#)

### 1 External memory algorithms and data structures: dealing with **massive** data 80%



Jeffrey Scott Vitter

**ACM Computing Surveys (CSUR)** June 2001

Volume 33 Issue 2

Data sets in large applications are often too massive to fit completely inside the computers internal memory. The resulting input/output communication (or I/O) between fast internal memory and slower external memory (such as disks) can be a major performance bottleneck. In this article we survey the state of the art in the design and analysis of external memory (or EM) algorithms and data structures, where the goal is to exploit locality in order to reduce the I/O costs. We consider a varie ...

### 2 A scalable distributed information management system 80%



Praveen Yalagandula , Mike Dahlin

**ACM SIGCOMM Computer Communication Review , Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications** August 2004

Volume 34 Issue 4

We present a Scalable Distributed Information Management System (SDIMS) that *aggregates* information about large-scale networked systems and that can serve as a basic building block for a broad range of large-scale distributed applications by providing detailed views of nearby information and summary views of global information. To serve as a basic building block, a SDIMS should have four properties: scalability to many nodes and attributes, flexibility to accommodate a broad range of appl ...

### 3 Data dissemination and pervasive computing: Power-efficient data dissemination in wireless sensor networks 77%




Ugur Cetintemel , Andrew Flinders , Ye Sun

**Proceedings of the 3rd ACM international workshop on Data engineering for wireless and mobile access** September 2003

This paper presents a new event-based communication model for wireless multi-hop networks of energy-constrained devices such as sensor networks. The network is arranged as an event dissemination tree, with nodes subscribing to the event types they are interested in. An event scheduler dynamically allocates and multiplexes upstream and downstream time slots for each event type. Power consumption among wireless nodes is reduced by allowing each node to power down its radio during the portions of t ...


**4 Fast detection of communication patterns in distributed executions** 77%

 Thomas Kunz , Michiel F. H. Seuren

**Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research** November 1997

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

**5 A case for dynamic view management** 77%

 Yannis Kotidis , Nick Roussopoulos

**ACM Transactions on Database Systems (TODS)** December 2001  
Volume 26 Issue 4

Materialized aggregate views represent a set of redundant entities in a data warehouse that are frequently used to accelerate On-Line Analytical Processing (OLAP). Due to the complex structure of the data warehouse and the different profiles of the users who submit queries, there is need for tools that will automate and ease the view selection and management processes. In this article we present DynaMat, a system that manages dynamic collections of materialized aggregate views in a data warehous ...

**6 Process migration** 77%

 **ACM Computing Surveys (CSUR)** September 2000

Volume 32 Issue 3

Process migration is the act of transferring a process between two machines. It enables dynamic load distribution, fault resilience, eased system administration, and data access locality. Despite these goals and ongoing research efforts, migration has not achieved widespread use. With the increasing deployment of distributed systems in general, and distributed operating systems in particular, process migration is again receiving more attention in both research and product development. As hi ...

**7 The holodeck ray cache: an interactive rendering system for global** 77%

 illumination in nondiffuse environments


Gregory Ward , Maryann Simmons

**ACM Transactions on Graphics (TOG)** October 1999

Volume 18 Issue 4

We present a new method for rendering complex environments using interactive, progressive, view-independent, parallel ray tracing. A four-dimensional holodeck data structure serves as a rendering target and caching mechanism for interactive walk-throughs of nondiffuse environments with full global illumination. Ray sample density varies locally according to need, and on-demand ray computation is supported in a parallel implementation. The holodeck file is stored on disk and ...



- 8** Workshop on compositional software architectures: workshop report 77%  
 **ACM SIGSOFT Software Engineering Notes** May 1998  
Volume 23 Issue 3

---

**Results 1 - 8 of 8      short listing**

---

The ACM Portal is published by the Association for Computing Machinery. Copyright ?2005 ACM, Inc.



[> home](#) [> about](#) [> feedback](#) [> login](#)

USPTO



Try the *new* Portal design

Give us your opinion after using it.

## Search Results

Search Results for: [ **topology <and> cache <and> evaluation <and> distributed <and> query <and> plan <and> directory <and> answer** ]  
Found **33** of **154,226** searched.

## Search within Results



[> Advanced Search](#)

[> Search Help/Tips](#)

Sort by: **Title** **Publication** **Publication Date** **Score** Binder

Results 1 - 20 of 33 short listing

Prev  
Page

1

2

Next  
Page

- 1** Per-user profile replication in mobile environments: algorithms, analysis, and simulation results 82%  
Narayanan Shivakumar , Jan Jannink , Jennifer Widom  
**Mobile Networks and Applications** October 1997  
Volume 2 Issue 2  
We consider per-user profile replication as a mechanism for faster location lookup of mobile users in a personal communications service system. We present a minimum-cost maximum-flow based algorithm to compute the set of sites at which a user profile should be replicated given known calling and user mobility patterns. We show the costs and benefits of our replication algorithm against previous location lookup approaches through analysis. We also simulate our algorithm against other location ...
- 2** Fast detection of communication patterns in distributed executions 82%  
Thomas Kunz , Michiel F. H. Seuren  
**Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research** November 1997  
Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...
- 3** A case for dynamic view management 80%  
Yannis Kotidis , Nick Roussopoulos  
**ACM Transactions on Database Systems (TODS)** December 2001  
Volume 26 Issue 4

Materialized aggregate views represent a set of redundant entities in a data warehouse that are frequently used to accelerate On-Line Analytical Processing (OLAP). Due to the complex structure of the data warehouse and the different profiles of the users who submit queries, there is need for tools that will automate and ease the view selection and management processes. In this article we present DynaMat, a system that manages dynamic collections of materialized aggregate views in a data warehouse ...

#### 4 External memory algorithms and data structures: dealing with **massive** data 80%



Jeffrey Scott Vitter

**ACM Computing Surveys (CSUR)** June 2001

Volume 33 Issue 2

Data sets in large applications are often too massive to fit completely inside the computers internal memory. The resulting input/output communication (or I/O) between fast internal memory and slower external memory (such as disks) can be a major performance bottleneck. In this article we survey the state of the art in the design and analysis of external memory (or EM) algorithms and data structures, where the goal is to exploit locality in order to reduce the I/O costs. We consider a varie ...

#### 5 Query evaluation techniques for large databases 80%



Goetz Graefe

**ACM Computing Surveys (CSUR)** June 1993

Volume 25 Issue 2

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

#### 6 A scalable distributed information management system 80%



Praveen Yalagandula , Mike Dahlin

**ACM SIGCOMM Computer Communication Review , Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications** August 2004

Volume 34 Issue 4

We present a Scalable Distributed Information Management System (SDIMS) that *aggregates* information about large-scale networked systems and that can serve as a basic building block for a broad range of large-scale distributed applications by providing detailed views of nearby information and summary views of global information. To serve as a basic building block, a SDIMS should have four properties: scalability to many nodes and attributes, flexibility to accommodate a broad range of appl ...

#### 7 Networking support: MEADOWS: modeling, emulation, and analysis of data of wireless sensor networks 77%



Qiong Luo , Lionel M. Ni , Bingsheng He , Hejun Wu , Wenwei Xue

**Proceedings of the 1st international workshop on Data management for sensor networks: in conjunction with VLDB 2004** August 2004

In this position paper, we present MEADOWS, a software framework that we are building at HKUST for modeling, emulation, and analysis of data of wireless sensor

networks. This project is motivated by the unique need of intertwining modeling, emulation, and data analysis in studying sensor databases. We describe our design of basic data analysis tools along with an initial case study on HKUST campus. We also report our progress on modeling power consumption for sensor databases and on wireless sen ...

## 8 Implementation and Evaluation of a Scalable Application-Level 77%



### Checkpoint-Recovery Scheme for MPI Programs

Martin Schulz , Greg Bronevetsky , Rohit Fernandes , Daniel Marques , Keshav Pingali , Paul Stodghill

**Proceedings of the 2004 ACM/IEEE conference on Supercomputing** November 2004

The running times of many computational science applications are much longer than the mean-time-to-failure of current high-performance computing platforms. To run to completion, such applications must tolerate hardware failures. Checkpoint-and-restart (CPR) is the most commonly used scheme for accomplishing this - the state of the computation is saved periodically on stable storage, and when a hardware failure is detected, the computation is restarted from the most recently saved state. Most aut ...

## 9 Interoperability of multiple autonomous databases 77%



Witold Litwin , Leo Mark , Nick Roussopoulos

**ACM Computing Surveys (CSUR)** September 1990  
Volume 22 Issue 3

Database systems were a solution to the problem of shared access to heterogeneous files created by multiple autonomous applications in a centralized environment. To make data usage easier, the files were replaced by a globally integrated database. To a large extent, the idea was successful, and many databases are now accessible through local and long-haul networks. Unavoidably, users now need shared access to multiple autonomous databases. The question is what the corresponding methodology ...

## 10 Technical papers: 4+4: an architecture for evolving the Internet address 77%



### space back toward transparency

Zoltán Turányi , András Valkó , Andrew T. Campbell

**ACM SIGCOMM Computer Communication Review** October 2003  
Volume 33 Issue 5

We propose 4+4, a simple address extension architecture for Internet that provides an evolutionary approach to extending the existing IPv4 address space in comparison to more complex and disruptive approaches best exemplified by IPv6 deployment. The 4+4 architecture leverages the existence of Network Address Translators (NATs) and private address realms, and importantly, enables the return to end-to-end address transparency as the incremental deployment of 4+4 progresses. During the transition t ...

## 11 Database session 8: interactive data exploration: Hierarchical graph 77%



### indexing

James Abello , Yannis Kotidis

**Proceedings of the twelfth international conference on Information and knowledge management** November 2003

Traffic analysis, in the context of Telecommunications or Internet and Web data, is crucial for large network operations. Data in such networks is often provided as large graphs with hundreds of millions of vertices and edges. We propose efficient techniques for managing such graphs at the storage level in order to facilitate its

processing at the interface level(visualization). The methods are based on a hierarchical decomposition of the graph edge set that is inherited from a hierarchical deco ...

## 12 Special topic section on peer to peer data management: DBGlobe: a service-oriented P2P system for global computing 77%



Evaggelia Pitoura , Serge Abiteboul , Dieter Pfoser , George Samaras , Michalis Vazirgiannis

**ACM SIGMOD Record** September 2003

Volume 32 Issue 3

The challenge of peer-to-peer computing goes beyond simple file sharing. In the DBGlobe project, we view the multitude of peers carrying data and services as a superdatabase. Our goal is to develop a data management system for modeling, indexing and querying data hosted by such massively distributed, autonomous and possibly mobile peers. We employ a service-oriented approach, in that data are encapsulated in services. Direct querying of data is also supported by an XML-based query language. In t ...

## 13 Peer-to-peer: Making gnutella-like P2P systems scalable 77%



Yatin Chawathe , Sylvia Ratnasamy , Lee Breslau , Nick Lanham , Scott Shenker

**Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications** August 2003

Napster pioneered the idea of peer-to-peer file sharing, and supported it with a centralized file search facility. Subsequent P2P systems like Gnutella adopted decentralized search algorithms. However, Gnutella's notoriously poor scaling led some to propose distributed hash table solutions to the wide-area file search problem. Contrary to that trend, we advocate retaining Gnutella's simplicity while proposing new mechanisms that greatly improve its scalability. Building upon prior research [1, 1 ...

## 14 Applications: Emergent properties of referral systems 77%



Pinar Yolum , Munindar P. Singh

**Proceedings of the second international joint conference on Autonomous agents and multiagent systems** July 2003

Agents must decide with whom to interact, which is nontrivial when no central directories are available. A classical decentralized approach is referral systems, where agents adaptively give referrals to one another. We study the emergent properties of referral systems, especially those dealing with their quality, efficiency, and structure. Our key findings are (1) pathological graph structures can emerge due to some neighbor selection policies and (2) if these are avoided, quality and efficiency ...






## 15 Distributed information retrieval: SETS: search enhanced by topic segmentation 77%



Mayank Bawa , Gurmeet Singh Manku , Prabhakar Raghavan

**Proceedings of the 26th annual international ACM SIGIR conference on Research and development in informaion retrieval** July 2003

We present SETS, an architecture for efficient *search* in peer-to-peer networks, building upon ideas drawn from machine learning and social network theory. The key idea is to arrange participating sites in a *topic-segmented* overlay topology in which most connections are *short-distance*, connecting pairs of sites with similar content. Topically focused sets of sites are then joined together into a single network by *long-distance* links. Queries are matched and ro ...

- 16** Dynamic services and analysis: Make it fresh, make it quick: searching a network of personal webserver 77%  
 Mayank Bawa , Roberto J. Bayardo , Sridhar Rajagopalan , Eugene J. Shekita  
**Proceedings of the twelfth international conference on World Wide Web** May 2003  
 Personal webserver have proven to be a popular means of sharing files and peer collaboration. Unfortunately, the transient availability and rapidly evolving content on such hosts render centralized, crawl-based search indices stale and incomplete. To address this problem, we propose YouSearch, a distributed search application for personal webserver operating within a shared context (e.g., a corporate intranet). With YouSearch, search results are always fast, fresh and complete -- properties we ...
- 17** Astrolabe: A robust and scalable technology for distributed system monitoring, management, and data mining 77%  
 Robbert Van Renesse , Kenneth P. Birman , Werner Vogels  
**ACM Transactions on Computer Systems (TOCS)** May 2003  
 Volume 21 Issue 2  
 Scalable management and self-organizational capabilities are emerging as central requirements for a generation of large-scale, highly dynamic, distributed applications. We have developed an entirely new distributed information management system called Astrolabe. Astrolabe collects large-scale system state, permitting rapid updates and providing on-the-fly attribute aggregation. This latter capability permits an application to locate a resource, and also offers a scalable way to track sys ...
- 18** IS '97: model curriculum and guidelines for undergraduate degree programs in information systems 77%  
 Gordon B. Davis , John T. Gorgone , J. Daniel Cougar , David L. Feinstein , Herbert E. Longenecker  
**ACM SIGMIS Database , Guidelines for undergraduate degree programs on Model curriculum and guidelines for undergraduate degree programs in information systems** December 1996  
 Volume 28 Issue 1
- 19** Papers: A survey of web caching schemes for the Internet 77%  
 Jia Wang  
**ACM SIGCOMM Computer Communication Review** October 1999  
 Volume 29 Issue 5  
 The World Wide Web can be considered as a large distributed information system that provides access to shared data objects. As one of the most popular applications currently running on the Internet, the World Wide Web is of an exponential growth in size, which results in network congestion and server overloading. Web caching has been recognized as one of the effective schemes to alleviate the service bottleneck and reduce the network traffic, thereby minimize the user access latency. In this pap ...
- 20** Dynamically distributed query evaluation 77%  
 Trevor Jim , Dan Suciu  
**Proceedings of the twentieth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems** May 2001

---

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.



[> home](#) [> about](#) [> feedback](#) [> login](#)

USPTO



Try the *new* Portal design

Give us your opinion after using it.

## Search Results

Search Results for: [ **topology <and> cache <and> evaluation <and> distributed <and> query <and> plan <and> directory <and> answer** ]

Found **33** of **154,226** searched.

## Search within Results



[> Advanced Search](#)

[> Search Help/Tips](#)

Sort by: **Title** **Publication** **Publication Date** **Score**  **Binder**

Results 21 - 33 of 33 short listing

   
Prev  
Page

1 2

   
Next  
Page

### 21 Process migration

77%



**ACM Computing Surveys (CSUR)** September 2000

Volume 32 Issue 3

Process migration is the act of transferring a process between two machines. It enables dynamic load distribution, fault resilience, eased system administration, and data access locality. Despite these goals and ongoing research efforts, migration has not achieved widespread use. With the increasing deployment of distributed systems in general, and distributed operating systems in particular, process migration is again receiving more attention in both research and product development. As hi ...

### 22 The proposed new Computing Reviews classification scheme

77%



Anthony Ralston

**Communications of the ACM** July 1981

Volume 24 Issue 7

### 23 The new (1982) Computing Reviews classification system—final version

77%



Jean E. Sammet , Anthony Ralston

**Communications of the ACM** January 1982

Volume 25 Issue 1

### 24 Session summaries from the 17th symposium on operating systems

77%



principle (SOSP'99)

Jay Lepreau , Eric Eide

**ACM SIGOPS Operating Systems Review** April 2000

Volume 34 Issue 2



**25** Workshop on compositional software architectures: workshop report 77%**ACM SIGSOFT Software Engineering Notes** May 1998

Volume 23 Issue 3

**26** Information gathering in the World-Wide Web: the W3QL query 77%

language and the W3QS system

David Konopnicki , Oded Shmueli

**ACM Transactions on Database Systems (TODS)** December 1998

Volume 23 Issue 4

The World Wide Web (WWW) is a fast growing global information resource. It contains an enormous amount of information and provides access to a variety of services. Since there is no central control and very few standards of information organization or service offering, searching for information and services is a widely recognized problem. To some degree this problem is solved by "search services," also known as "indexers," such as Lycos, AltaVista, Yahoo, and others. ...

**27** Analyzing stability in wide-area network performance 77%

Hari Balakrishnan , Mark Stemm , Srinivasan Seshan , Randy H. Katz

**ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1997 ACM SIGMETRICS international conference on Measurement and modeling of computer systems** June 1997

Volume 25 Issue 1

The Internet is a very large scale, complex, dynamical system that is hard to model and analyze. In this paper, we develop and analyze statistical models for the observed end-to-end network performance based on extensive packet-level traces (consisting of approximately 1.5 billion packets) collected from the primary Web site for the Atlanta Summer Olympic Games in 1996. We find that observed mean throughputs for these transfers measured over 60 million complete connections vary widely as a funct ...

**28** Technique for automatically correcting words in text 77%

Karen Kukich

**ACM Computing Surveys (CSUR)** December 1992

Volume 24 Issue 4

Research aimed at correcting words in text has focused on three progressively more difficult problems: (1) nonword error detection; (2) isolated-word error correction; and (3) context-dependent word correction. In response to the first problem, efficient pattern-matching and n-gram analysis techniques have been developed for detecting strings that do not appear in a given word list. In response to the second problem, a variety of general and application-specific spelling cor ...

**29** A taxonomy of issues in name systems design and implementation 77%

A. K. Yeo , A. L. Ananda , E. K. Koh

**ACM SIGOPS Operating Systems Review** July 1993

Volume 27 Issue 3

In the last decade, name systems have grown from a single centrally-controlled server providing only host name to physical address mapping, to a complex system consisting of multiple and distributed servers, providing not only name mapping, but also general directory lookup services. These advances are due in part to the increase in size, complexity and heterogeneity of distributed systems. This paper presents a taxonomy of design and implementation issues in building a name system.

**30** Distributed indexing: a scalable mechanism for distributed information retrieval 77%

Peter B. Danzig , Jongsuk Ahn , John Noll , Katia Obraczka

**Proceedings of the 14th annual international ACM SIGIR conference on Research and development in information retrieval** September 1991**31** Distributed data sources: Efficient query routing in distributed spatial databases 77%

Roger Zimmermann , Wei-Shinn Ku , Wei-Cheng Chu

**Proceedings of the 12th annual ACM international workshop on Geographic information systems** November 2004

Spatial databases are prominently used in Geographic Information System (GIS) applications. However, many of the current architectures rely on a centralized data repository. The next evolution will be GIS applications that utilize and integrate a multitude of remotely accessible data sets, for example via Web services. Our involvement in a project where geotechnical borehole information is retrieved from a large number of repositories that are under different administrative control has motiva ...

**32** Service discovery in agent-based pervasive computing environments 77%

Olga Ratsimor , Dipanjan Chakraborty , Anupam Joshi , Timothy Finin , Yelena Yesha

**Mobile Networks and Applications** December 2004

Volume 9 Issue 6

Directory based service discovery mechanisms are unsuitable for ad-hoc m-commerce environments. Working towards finding an alternate mechanism, we developed Allia: a peer-to-peer caching based and policy-driven agent-service discovery framework that facilitates cross-platform service discovery in ad-hoc environments. Our approach achieves a high degree of flexibility in adapting itself to changes in ad-hoc environments and is devoid of common problems associated with structured compound forma ...

**33** System support for pervasive applications 77%

Robert Grimm , Janet Davis , Eric Lemar , Adam Macbeth , Steven Swanson , Thomas Anderson , Brian Bershad , Gaetano Borriello , Steven Gribble , David Wetherall

**ACM Transactions on Computer Systems (TOCS)** November 2004

Volume 22 Issue 4

Pervasive computing provides an attractive vision for the future of computing. Computational power will be available everywhere. Mobile and stationary devices will dynamically connect and coordinate to seamlessly help people in accomplishing their tasks. For this vision to become a reality, developers must build applications that constantly adapt to a highly dynamic computing environment. To make the developers' task feasible, we present a system architecture for pervasive computing, called & ...

**Results 21 - 33 of 33** short listing
  
 Prev  
 Page

1 2

  
 Next  
 Page

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.